CHAPTER 2

CONTROL SYSTEM NETWORK

2-1 INTRODUCTION

This chapter describes the building-level open-communications control system architecture and control devices for HVAC and other building-level monitoring and control applications. The communications network and devices are based on LonWorks® technology and ANSI-709.1 communications protocol.

Design of an open-communications building-level control system does not require an extensive familiarity with the ANSI-709.1 protocol. Pertinent information in regard to control system design is presented in this document. System architecture, device functionality and data sharing capabilities are also described in this chapter.

2-2 ADDRESSING, DATA TRANSMISSION, AND DATA INTEGRITY

All network protocols (such as ANSI-709.1) define an addressing scheme: a method of delivering messages to a specific device on the network. ANSI-709.1 defines several such methods, the two most common ones are:

- Information can be sent using the recipient's unique NodeID (more commonly referred to as the Neuron ID) as the destination address. This address, which is a 48 bit number, is set at the factory for every ANSI-709.1 device and is guaranteed to be unique – no two ANSI-709.1 devices will ever have the same NodeID.
- Information can be sent using a destination address of the form Domain/Subnet/Node. The installer will assign addresses to DDC Hardware according to base-specific UMCS-wide guidelines. The Domain is the top-level portion of the address and is generally fixed; most installations will have one Domain address for the whole UMCS (except for large installations). The subnet address will generally correspond to a single building, though larger buildings may require several subnet addresses while multiple small buildings may share a common subnet address. At the lowest level are individual node addresses; each DDC Hardware on a subnet must have a unique node address.

  This method requires more care to set up than the use of NodeIDs. The big advantage of this method is that the addressing scheme can/should reflect the logical organization of the control network. For example, a large building may consist of subnets 105 – 110, with a large AHU controller at subnet 108, device 1 and its associated VAV boxes at subnet 108, devices 2-35 (whereas if NodeIDs were used, the addresses would all essentially be random numbers between 1 and 281 trillion).

The specification requires the use of the Domain/Subnet/Node addressing scheme with documentation of the NodeIDs for DDC Hardware.

While the underlying wire speed is 78.1 kbps, there are several aspects of the protocol that govern how much bandwidth is used by a given data transfer. There are 2 main methods to transfer data within ANSI-709.1, polling and binding. Polling occurs when a receiver of data requests data from a transmitter; this is generally a periodic event with a well-defined period. A classic example of polling is trend data from the UMCS; every

15 minutes the front-end will request data from several DDC hardware devices. Polling can occur at any time; a device can always poll another device for data.

The other form of data transfer, binding, is set up at network configuration time and generally occurs when a data source sends (on its own initiative) data to a recipient. There are several parameters that control the frequency of this transfer:

- Change of Value. The transmitter can be configured to only send the data if it changes by a minimum specified amount. For example, a OA-T sensor might be configured to send a new temperature value only if the current value changes from the last sent value by greater than .5 degrees.
- Minimum send time. The transmitter can be configured not to send the data more often than once every ***X*** seconds. This is an important parameter for limiting network traffic; most HVAC control applications (except duct static pressure control or flow matching) does not require "real-time" data and data should probably not be sent more often than once every couple of seconds even if it is rapidly changing.
- Maximum send time. The transmitter can be configured to send the data at least once every ***X*** seconds, even if the value is not changing. This is generally a good idea just in case something goes wrong. For example, if the receiving device is reset, it won't "remember" the data and will not have the value until it is resent. A typical maximum send time might be 30 minutes.

Finally, there are several other parameters that govern data integrity: how the protocol ensures reliable data communication with less-than-perfect hardware, noisy lines, "glitches" and other real-world events. With binding, there are several common means to send data, each has advantages and disadvantages:

- Unacknowledged send once. The data is sent one time and one time only. This requires the least network bandwidth, but does not provide any assurance that the data reaches the recipient. This is the most common form of data binding.
- Unacknowledged send multiple. The data is sent multiple times (typically 3). It is up to the recipient to deal with receiving the same data multiple times. This requires more bandwidth than the send once option, but is still fairly fast because the transmitter does not wait for any acknowledgement, it simply sends the data ***X*** times and moves on.
- Acknowledged send. The data is sent once. Upon receipt of the data, the receiver must send an acknowledgement message back to the transmitter. If the transmitter does not receive the acknowledgment within a pre-determined period of time (the "timeout"), the transmitter will resend the data. This is the slowest method, but is a good trade-off between reliability and network bandwidth usage. This is typically used for alarms, where it is essential that the data get transferred.

2-3 NETWORK BANDWIDTH CALCULATION

The specification requires that the building contractor perform two network bandwidth calculations to ensure that the proposed system (devices, network bindings, and network architecture) does not saturate the network, one for a heavily loaded network and one for a normally loaded network. For each case, the specification defines a collection of network activities that will consume network bandwidth. For each activity, the contractor must consider the bandwidth used by that activity, which will factor in the number of SNVTs, the network topology, the data transmission rate (which will depend

on the applications) and the data integrity method used for sending the data (which will also depend on the applications).

2-4 **MEDIA**

2-4.1 **TP/FT-10 Media**

Selection of the physical media type is an important part of specifiying a DDC system. Media type will largely determine network speed (bandwidth) and will set a maximum length for a single cable connecting multiple DDC Hardware. These cable length restrictions are due to simple signal degradation and also due to timing issues associated with fast network signals on longer cables.

Buildings that cannot be served by a single cable will need to be broken into multiple sub networks, which will be connected with routers and/or repeaters into a larger building-wide network. Note that there other reasons (explained below) why it is advantageous to utilize multiple sub networks in a building.

The building level control network will use the ANSI-709.1 communications protocol (ANSI/EIA/CEA-709.1B) over a TP/FT-10 network connected in a doubly-terminated topology. The term “TP/FT-10” defines a network media and transceiver type:

- The TP in TP/FT-10 stands for Twisted Pair. This is a description of the media that is used to connect the controllers. In this case, a twisted pair of wires is used. ANSI 709.3 requires that this twisted pair meets the requirements of CAT-5 cable. While the protocol will work over a variety of cable types, CAT-5 (or better) cable is such a widely used standard that requiring the use of it will help avoid incompatibility problems later.
- The FT in TP/FT-10 stands for Free Topology and indicates the transceiver type that controllers on the network will use. The transceiver is responsible for actually transmitting information across the network (across the media). Note that while this allows for Free Topology, the specification further restricts the network to a doubly-terminated bus topology.

Doubly-Terminated Bus Topology requires that the bus be daisy-chained from one device to another with no branches (stubs under 3 meters in length are allowed in accordance with ANSI 709.3) with terminators at both ends of the bus. The spec requires doubly-terminated bus topology in order to maintain consistency and since this topology is the easiest to understand/work with.

2-4.2 **IP Network Media**

The building level control network may also include an IP network if required as discussed below, as specified in Section 13801 Utility Monitoring and Control Systems and as described in UFC 3-401-01.

2-4.3 **Other Media Types**

In addition to TP/FT-10, there are two media/network types that are part of the ANSI-709 standard, Power Line (ANSI-709.2) and Fiber Optic (ANSI-709.4). Furthermore, there are many media/network types available that are not included in the ANSI

standard. Many of these media types should be avoided, but some (most notably Radio Frequency (RF)) may be useful in some applications. TP/FT-10 is the default and recommended media type since it is the most supported and open option. Use of other media types may limit future competition by giving an advantage to the limited number of vendors whose products support the non-standard media.

The determination of whether to allow/specify alternative media types is best made by asking “What am I gaining by using this media instead of TP/FT-10?” and “What am I losing by using this media instead of TP/FT-10?”. Often the answer to the first question will be that it is a matter of convenience, while the answer to the second will be that the system will become less open. In these cases, it is often worthwhile to proceed with TP/FT-10 despite the additional cost/time, as it will prove to be more convenient in the long term.

In general, if the alternative media type calls for the installation of media, then there is little or no benefit to using that alternative media. If the alternative media allows the use of existing media (Power Line, RF or Fiber Optic), there may be motivation to allow it, but the impact on the openness of the system must be considered.

Since TP/FT-10 offers the most compatible devices, alternative media (with the possible exception of PL) should only be allowed when it is used in conjunction with TP/FT-10:

- To bridge 2 TP/FT-10 segments
- As a local control bus connected to a TP/FT-10 backbone.

2-5 **NETWORK HARDWARE**

In addition to media, the control network may contain the following types of hardware.

2-5.1 **Repeater**

A repeater is a device that has 2 or more input/output ports, connects two (or more) pieces of media, and performs signal regeneration. Signals showing up on an input port get cleaned up, amplified, and sent out of the repeaters output port(s). Repeaters may allow for longer cable runs in some cases, but not others.

2-5.2 **Media Converter**

A Media Converter is a repeater that changes media types, i.e. TP/FT-10 to PL. Use of non-standard media will probably require the use of media converters at each end of the non-standard media.

2-5.3 **ANSI-709.1 Router**

A router is similar to a repeater, but performs the additional function of packet filtering based on destination address. A router will look at the destination address of an incoming packet. If the destination DDC Hardware is accessible via media connected to a different output port, the packet will be sent out the appropriate output, otherwise the router will do nothing with the packet.

Routers maintain routing tables which list which domains and subnets exist on its output ports. Routers typically will also contain a “default” entry, which essentially says “If the destination doesn't show up in any routing table entry, forward the packet to another (specified) device (and hope that device can forward it properly).”

Routers may be learning or configured. A configured router has its routing tables assigned by the installer. A learning router will “learn” its routing tables. Initially, a learning router simply functions as a repeater and forwards all messages. As messages pass through the router, it looks at the source subnet address and learns which of its input ports connects to that subnet; it can then use that information to build a routing table entry for that subnet.

Routers serve two very important functions in a control network:

- First it greatly reduces network traffic. By placing devices that need to communicate frequently on a common subnet and isolating that section with a router, the base-wide network will not be bogged down with local communications between those devices.
- Second, it allows devices to send messages to a “distant” controller without knowing the detailed network topology. A device (say an OA-T sensor) in one building that wishes to send the temperature to another device in another building only needs to forward the message to its router; the router is then responsible for knowing how to send the message on towards the destination device.

2-6 **ARCHITECTURE**

Figure 2-1 shows a sample architecture for a basewide system consisting of multiple building level controls networks (installed per UFGS-15951 and this UFC) connected to a Utility Monitoring and Control System (UMCS).

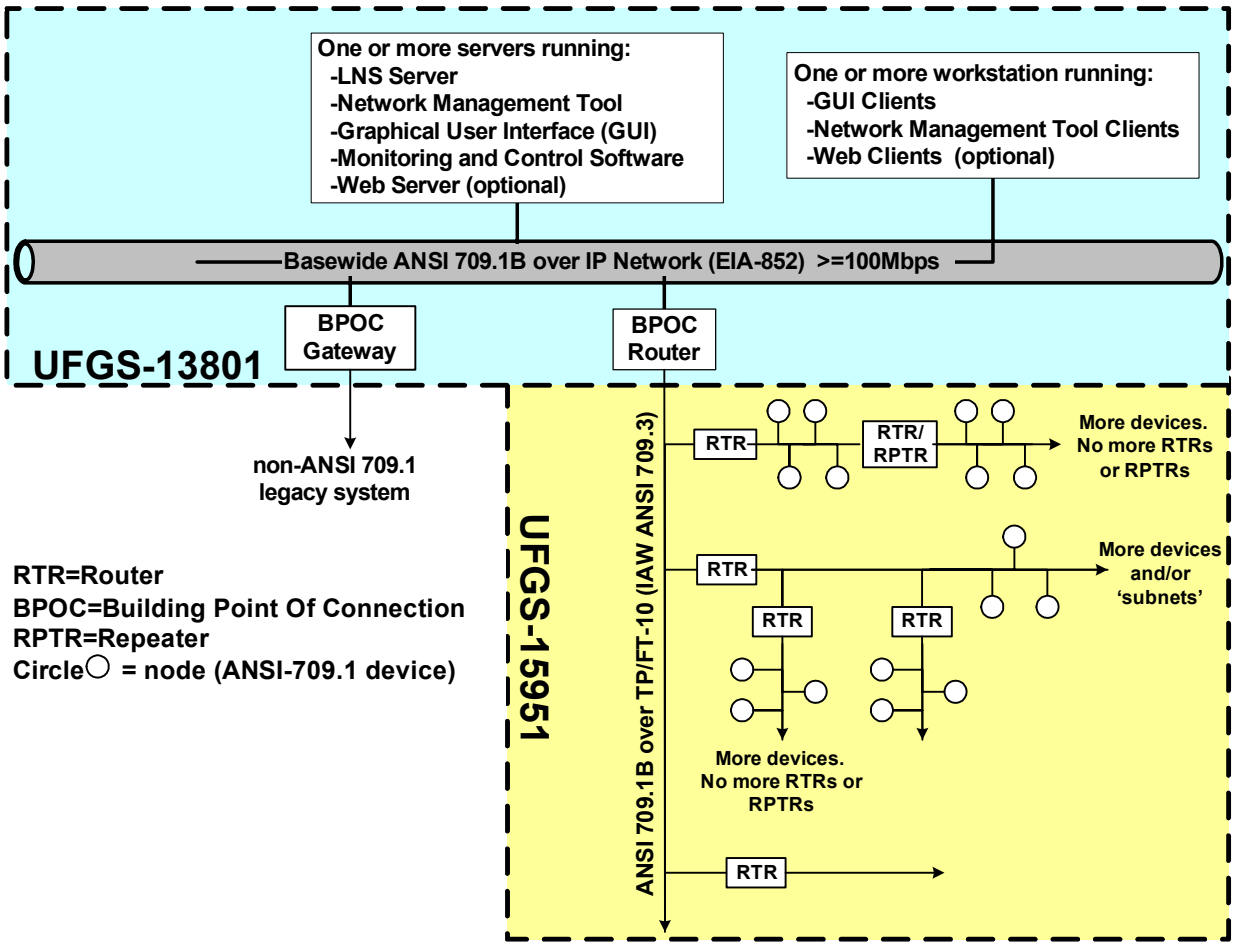

Figure 2-1: Sample Architecture

The building level control network will use the ANSI-709.1 communications protocol (ANSI/EIA/CEA-709.1B) over a TP/FT-10 network in bus topology. This network will consist of a *backbone* with one or more *local control buses* connected to it using routers. This produces a logically flat network in the building where each node can communicate directly with any other node without the intervention of another controller. Some possible variations to this basic architecture are:

- Multiple buildings can share a common building-level backbone. For example, two or more adjacent buildings can be physically linked by a common 78 Kbps backbone network cable as long as network restrictions such as cable length and the total number of nodes as described elsewhere in this document and in UFGS-15951 are adhered to. If connecting the building to a UMCS in this case a single BPOC can be used to connect these buildings to the UMCS.

- A large building can have multiple building-level backbones. In this case, the building-level contractor(s) will install multiple independent building-level ANSI-709.1 networks. If connecting the building to a UMCS in this case the UMCS contractor will install a separate BPOC for each network

2-6.1 **Local Control Buses**

The Local Control Bus is the portion of the building network where controllers are connected. The guide specification requires that the local control bus be a TP/FT-10 network (in double-terminated bus) topology as specified in ANSI 709.3.

2-6.2 **Building Network Backbone**

The building network backbone connects the different local control buses in the building together and connects to the UMCS network. The specifications require that the only devices on the backbone be routers.

With some exceptions, the backbone will be a TP/FT-10 network in doubly-terminated bus topology. However, the contractor is required to perform network bandwidth usage calculations on the backbone to determine if the TP/FT-10 network will prove too slow in application. If these calculations show a bandwidth usage of over 70% of the 78 kbps available the contractor is required to instead install an IP network. Since it is expected that in most/all cases the TP/FT-10 network will meet the requirements, IP network installation requirements are not included in Section 15951, and the spec instead requires that the IP network be installed per Section 13801.

2-6.3 **Additional Architecture issues**

2-6.3.1 **Local Control bus vs Backbone Bus**

Segregating the network into local control buses and a backbone is the easiest way to manage network traffic and not overload the network. The intent is that devices that need to communicate frequently be placed on a common local control bus. The requirement that only routers be connected to the backbone ensures that traffic from a (potentially) congested local control bus does not clog the backbone – the router will keep local traffic on the local control bus and off the backbone. Traffic between devices that communicate infrequently (or inter-building communication, such as with the front-end UMCS) will utilize the building backbone. The requirement that no node has more than 2 routers/repeaters between it and the backbone helps ensure that the installer doesn't bog down a local control bus by forcing traffic from a second local control bus to traverse the first local bus to get to the backbone.

2-6.3.2 **Need for IP backbones**

While a networked DDC system offers enormous advantages over a non-networked system, the actual requirements for data transfer over the network are relatively minor. Attention to several details will help ensure that the relatively slow 78 kbps TP/FT-10 network will suffice for the building backbone:

- Group devices that need to communicate often on a common local control bus
- Limit the amount of information sent to the UMCS. A modern UMCS can easily demand data from the local controls faster than the building network can deliver the data. Coordinate with the UMCS installer to limit “always-active” data requests from the UMCS such as trending to those really required by the installation.
- Ensure the contractor utilizes care in selecting data transfer rates and integrity methods. Use “Send on Change” with reasonable change values to avoid sending

data more often than required. Limit “Unacknowledged Send Multiple” and “Send Acknowledged” transmissions to critical data only.

2-6.3.3 **Multiple controllers per HVAC system vs single controller**

The LonWorks industry supports the idea of “distributed control”. In a conventional DDC system, a single relatively powerful controller will be used to control an AHU. While there are ANSI-709.1 controllers that support this option, another possibility is the use of multiple relatively simple interconnected controllers. In this case, the mixed air dampers may be a dedicated controller (or even a so-called smart actuator) which lives on the ANSI-709.1 control network, obtains relevant temperatures from other “smart” sensors also on the network, obtains occupancy status and other information from other devices on the network and drives the dampers. Similarly, the cooling coil valve would be driven by a simple controller whose only output would be a 4-20 mA control signal to the valve. (This approach is somewhat similar to the old SLDC approach, where an AHU controller would be built-up from multiple simple PI controllers).

While this approach to HVAC control is not necessarily unique to ANSI-709.1 based hardware, this approach does seem to be better supported by LonWorks than other protocol technologies. As with any approach, distributed control has its own set of advantages and disadvantages.

Advantages:

- Simple controllers: Programmable controllers not required, which eliminates custom programming and programming software. A small selection of simple controllers can be used for all control schemes. This would allow an installation to standardize on a set of controllers and ease the training requirements for the O&M Staff.
- Can be duplicated by multiple vendors: Since the devices have simple functionality it is easier to find a replacement device from a different vendor with the same functionality.
- Easier to document the actual controller and controller functions/settings.
- No long home-runs of wire. Controllers may be located at the sensors/actuators they interface, or the sensors/actuators may be the controllers (‘smart’ sensors/actuators) which reduced the wiring requirements.

Disadvantages:

- Execution of the system sequence may require communication between multiple controllers which in turn may require a functional network. One work-around to this issue is to create a local control network dedicated to the system and isolate this from the rest of the building with a router (this will protect the local network from most network failures elsewhere in the building).
- Harder to document. While the individual controllers were easier to document, the system sequence may be harder to document since it is distributed among multiple controllers.
- Controllers not in one location. The controllers for a single system may be scattered about the mechanical room, or even outside the mechanical room.

While it is preferred to avoid distributed control, there is no reason not to use it when and where it makes sense to do so. The Guide Spec places the burden on the contractor to decide when distributed control should be used.

2-7 CONNECTION TO A UMCS

As previously discussed, the building-level control network will perform all necessary control functionality in a stand-alone mode but does not provide an operator interface for monitoring and control of the network. If the building is to be operated in a stand-alone mode for an extended period and monitoring and control functionality are required, the designer should utilize the required portions of UFGS-13801 to obtain a local monitoring and control system. If the building is to be connected to the UMCS, the UMCS contractor will be responsible for installation and configuration of the BPOC and integration of the building system into the UMCS. See Chapter 6: Project Implementation for more information.

2-8 NETWORK DESIGN AND LAYOUT

Network layout is left largely to the building-level controls contractor as specified in UFGS-15951. Designer layout selections and considerations are described in the Project Implementation chapter.